

RECONSTRUCTION OF SAMPLED AND MODELED OCEANOGRAPHIC INFORMATION USING AN EXPLORATORY ENVIRONMENT

Peter Gruzinskas
David Cole
Naval Oceanographic Office
Stennis Space Center, Mississippi 39522

Andy Haas
Logicon, Inc.
Stennis Space Center, Mississippi 39522

Abstract

The Naval Oceanographic Office (NAVOCEANO) Department of Defense (DoD) Major Shared Resource Center (MSRC) operates as a DoD shared high performance computing (HPC) center, serving over 2000 users nationwide. Scientific visualization is an essential element of HPC that provides the methodology to explore, define, and present the results of computations. Accurate and understandable representation of data impacts a project in several ways: it helps scientists understand the physical phenomena they are studying, it helps communicate this work to their colleagues, and it helps explain the work and its significance to the public. Our scientific visualization team, in conjunction with researchers at our facility and other DoD facilities, have accomplished numerous visualizations of large data sets. In this paper we present examples of the methods, techniques, and results of our efforts to visualize observational data and the output of numeric ocean models.

I. Background

The staff at the Naval Oceanographic Office (NAVOCEANO) Major Shared Resource Center (MSRC) Visualization Lab is tasked with supporting the DoD research community by providing state-of-the-art scientific visualization support involving several Computational Technology Areas (CTA). Our MSRC provides tailored support for Computation Fluid Dynamics, Signal Image Processing, Climate Weather and Ocean Modeling, Environmental Quality Modeling, Computational Electromagnetics and Acoustics. We also work closely with various components of NAVOCEANO's operational community, to assist in the development of new techniques to display, analyze, and quality control some of the measurements and model output which provide critical operational support to the Navy. The common problem that this support must overcome is dealing with large data sets. "Large" is a fluid definition; it changes as technology evolves. In

today's environment we are dealing with datasets that range routinely from hundreds of megabytes to several gigabytes. Service to remote users also poses a challenge to our scientific visualization staff. Remote could be considered anywhere from a room in another building, to another building, in a different state, in a different time zone. Our user base is nationwide.

II. Introduction

The premise of this paper and presentation is to demonstrate the utility of an interactive digital environment to analyze oceanographic data and measurements. In such an environment one can gain insight into physical processes and ensure the quality of both measurements and model output by exploring and interrogating the data in both time and space. In this case we are dealing with environmental processes which are historically difficult to model and measure. Some if not all of these processes have impact on naval weapons and sensors. Our ability to understand, model, predict, and present these phenomena is important to our military. In our study we will analyze ocean bathymetry, circulation, and other ocean climatology over several domains. We will explore both space and time as we interrogate these data sets in 3-D computer space. Each data type has its own unique requirements and must be available to the display routine. The routine involves storage, networking, processing power, memory, and more to achieve this environment. Fortunately, these requirements are becoming readily available at low cost.

Modern oceanographic survey and modeling techniques, coupled with the successes of high performance computing, have begun to overwhelm both scientists and researchers with numerical information. The challenge of scientific visualization is to present the critical, distilled information to a varied client base, but in all cases, maintain the integrity of the data. These data appear to be increasing in resolution at an enormous

pace, both spatially and temporally. The term “large data set” will be continuously redefined.

Our 3-D environment will represent static measured boundary conditions, such as bathymetry, fused with dynamic modeled oceanographic parameters, such as circulation and temperature.

III. The Data

A. Measured

Digital elevation data obtained from a variety of public domain sources form the frame of reference in which to present our time series oceanographic information. It is critical when presenting this type of information to accurately define the surface and geophysical boundary conditions of the modeled environment. For the most part we will define our geophysical boundary with ETOPO5, but have enhanced some enclosed basins with higher resolution bathymetry. Therefore, our topography represents nominally 5-nmi horizontal resolution and our bathymetry anywhere from 5- to 1-nmi resolution. The techniques developed using these grid resolutions can and will be applied to higher resolution data. Work with different components within NAVOCEANO has involved viewing a variety of data types within the context of our terrain visualization. These data types include, but are not limited to, acoustic imagery, bioluminescence, circulation, salinity, sonic layer depth, and temperature.

B. Modeled

The modeled data that we present illustrates 3-D ocean circulation and create a large overhead in storage, memory, and input/output (I/O). Model output comes in a variety of formats, with varying resolutions in both the spatial and temporal domains. Determining the optimal sample rate and display resolution becomes the challenge of scientific visualization. Higher resolutions have larger storage/bandwidth requirements. One timestep of the Miami Isopycnal Coordinate Ocean Model (MICOM), which covers the Atlantic with 16 vertical layers, is approximately ½ GB.

The following describes the modeled output we visualize:

Numerical modelers at the University of Miami have long pursued the goal of studying the ocean circulation using models formulated in density (isopycnal). Because many physical processes in the ocean are rather intimately related to isopycnal surfaces and to the way in which they deviate from the horizontal, isopycnal models must have to contribute in elucidating oceanic circulation features with scales ranging from frontal to global. The association of vertical shear with isopycnal packing and tilting in the ocean makes these models appropriate for studies of strong baroclinic currents, such as the Gulf Stream. However, the fundamental reason for modeling

ocean flow in density coordinates is that this system suppresses the “diapycnal” component of numerically caused dispersion of material and thermodynamic properties. It is this characteristic that allows isopycnal models to keep deep water masses near the freezing level for centuries—in agreement with observation—while surface waters can be as warm as 30 degree Celsius. Models framed in Cartesian coordinates suffer from vertical “heat leakage” which cases the ocean to act as a giant heat sink in climate simulations.

The long-term goals of the modelers at the University of Miami are to perform a realistic, truly eddy-resolving wind- and buoyancy-forced numerical simulation of the North Atlantic Basin with data assimilation capabilities and to assess the nowcast/forecast capabilities of such a high-resolution ocean model. One of the primary research objectives is real-time forecasting of both Lagrangian trajectories and 3-D Eulerian fields associated with such physical parameters as velocity, temperature, salinity, and density. The five major components of the effort will be (1) MICOM, (2) data, (3) an Extended Kalman Filter (EKF) with a Gauss-Markov Random Field (GMRF) model for spatial covariances, (4) a random flight turbulence model for Lagrangian trajectory prediction, and (5) contour-based parameter estimation and assimilation techniques. The computational requirements for basin-scale ocean modeling at the resolutions of interest (less than 10 km) are extreme. Each time that the horizontal resolution is increased by a factor of n , the computational load goes up by a factor n^3 since the n -fold reduction in linear mesh size requires n times more time steps to integrate the model over a given time interval.

The main scientific goal will be to generate optimal estimates of the time-varying ocean state in support of the Navy’s needs on synoptic time scales on the order of weeks to months and on spatial scales typically on the order of 100 to 1000 km (mesoscale). Doing this in real time requires interplay between large varied data sets, numerical ocean circulation models, and data assimilation algorithms. Due to the large demand placed by near-optimal assimilation techniques on raw computing power, this work will fit most naturally under the Grand Challenge label.

The project’s fine-mesh simulation of the North Atlantic circulation (mesh size 0.08 degree longitude, 6 km on the average) is presently in its eighth year of integration. This simulation has generated considerable interest in both the computing and the oceanographic community. The term “fine-mesh” describes a horizontal grid resolution—typically of order 10 km—which allows barotropic/baroclinic instability, shear instabilities typical for geophysical flows, to be modeled. Since these instabilities cause ocean currents to meander and break up into individual eddies, fine-mesh ocean models are also referred to as “eddy-resolving.”

In this configuration, a realistic result for the Gulf Stream separation is achieved. This result supports the view that an inertial boundary layer (which results from the fine resolution) is an important factor in the separation process. It was also the first simulation with a fully thermodynamic basin-scale model to simulate the separation in a realistic fashion. This simulation allows direct and detailed comparisons with observations such as satellite data (sea surface height and sea surface temperature), mooring measurements, inverted echo sounders, and free-floating drogues.

The computational domain is the North and Equatorial Atlantic Ocean Basin from 28°S to 65°N (including the Caribbean Sea and the Gulf of Mexico but excluding the Mediterranean Sea) with a horizontal resolution of 1/12 degree (mesh size on the order of 6 km) and 16 layers in the vertical. The vertical grid was chosen to provide maximum resolution in the upper part of the ocean.

NAVOCEANO has developed an operational capability to forecast ocean currents and thermal structure in semienclosed seas. The Persian Gulf is one of the areas in which a numerical modeling system has been employed. The principal elements of the system are (1) a 3-D primitive equation circulation model, (2) temperature, salinity, and bathymetry data bases, and (3) meteorological forecasts provided by the Fleet Numerical Meteorology and Oceanography Center.

The model uses terrain-following vertical coordinates, where each level is a fixed fraction of the water depth. The spacing of the levels is reduced near the surface and bottom so that top and bottom boundary layers are resolved. The Persian Gulf models uses an along-axis resolution of 4.4 nmi and a cross-axis resolution of 4.6 nmi.

The Persian Gulf is a shallow, semienclosed basin with a mean depth of only 25 to 40 m. The circulation of this basin is driven primarily by the local wind stress and secondarily by thermohaline forcing. The prevailing wind in the Persian Gulf is from the northwest and is called the shamal. A wind-driven generally cyclonic circulation results. The lands surrounding the Persian Gulf are dry so there is strong excess evaporation over the Persian Gulf. This results in a surface inflow of relatively fresh water and an outflow of deeper, more-saline water at the Strait of Hormuz.

Model output shows the generally cyclonic circulation but with more complexity. Some of the highest current speeds are in the inflow through the southern side of the Strait of Hormuz. This inflow feeds the eastward coastal current along the south edge of the Gulf, which is strongest near Qatar. Along the Iranian coast, there is

another eastward current where it terminates and its remnant turns south into the interior.

IV. The Application

There are several generations of explorers built to view uniquely different data sets. Our strategy is to build the application around the data and provide low-cost portable tools which are tuned to the user's data visualization needs. In an interactive computer environment such as the one we describe, one must be able to deal with the geometry which describes the ocean basin before attempting to inject and interact with additional information, such as time-series circulation, temperature, or salinity. Rendering the terrain dynamically is accomplished using Dynamic Surface Generation (DSG). DSG uses layers of varying resolution (mip mapped) to provide optimum resolution vs. interactive performance with the option to view full resolution at any time. We plan to provide NAVOCEANO's Seafloor Data Bases Division a version of this software which will provide them a low-cost, portable means to assist in the exploration and quality control of their gridded data bases. They will push the envelope of our application, applying it to their high-resolution grids, which are nominally .1 arc minutes or approximately 180 m.

DSG is a simplification of the Silicon Graphics, Inc. (SGI) Active Surface Definition (ASD) technique. ASD maximizes terrain quality (both surface texture and geometry) while maintaining high frame-rate interactivity. It is intended for visual/flight simulation and supports real-time texturing of large landscapes. Its ability to roam large textures is based on clipmapping hardware available only on SGI Infinite Reality systems. To maintain portability, DSG does not use clipmapping. DSG builds surface geometry on the fly by computing the extent of longitude and latitude within the user's field of view. Multiple-resolution copies of the bathymetry data exist in a mipmap, which can be thought of as an inverted pyramid with the highest resolution layer on top and progressively smaller resolution layers working downward. The best resolution layer is chosen such that the number of polygons needed to cover the extent of longitudes and latitudes is kept to a minimum. This minimum can be set lower or higher to increase interactivity or resolution, respectively.

DSG is optimal when the view direction is perpendicular (i.e., looking top down) to the surface. As the view direction becomes parallel with the surface, the viewing area increases and drives DSG to use a coarser resolution to cover the increased area of data. Fly-through visual simulations, by their very nature, operate with view direction parallel to the surface. Because of this, ASD optimizes surface resolution by combining high-resolution data closest to the user with lower resolution data farther away. This is an involved process

that contrasts DSG's approach of maintaining a uniform resolution for all data in the field of view. This simpler approach works well in oceanographic applications where it is common to look at the data in a top-down fashion.

The viewing area used by the DSG is not only important for interactive terrain rendering, but for visualizing large data sets. Because the full range of data is often much larger than the memory and graphics can handle, the viewing area limits the amount of data that needs to be read. Data are read at a resolution that matches the terrain. When data are read across a large view area, a coarser version is produced to fit the memory and graphics. As the view area becomes smaller, the application can visualize the data closer to its original resolution. The application interleaves data I/O with its animation and user input. This allows data covering different areas/resolutions to be processed, while allowing the user to interact with the current data.

The data bounded by the topography are dense and curvilinear along one or more dimensions. The data contain velocity information that is visualized via two techniques: glyphs and particles. Glyphs are 3-D arrows that provide the viewer with a general feel of the vector data. They are constructed as cones with a tube body and rendered with OpenGL, which is the standard graphics library for 3-D rendering on PC's and workstations. The length and color of the glyphs map the ocean current speed. The color mapping defaults to speed, but can be mapped to other scalar information throughout the grid. The standard blue, green, red is used, but any table of colors can be supplied. The glyphs can be sampled at full or partial resolution across the data. Even at full resolution, glyphs by themselves are only clues about the flow of information hidden within the data. Particle advection is used to reconstruct the continuous sense of flow by interpolating velocities and grid layer depths (for isopycnal models) in space and time. Integrating particle velocities creates paths lines that are persistent across time. This persistence allows time-varying data to be studied more effectively than glyph animations. Path line colors are mapped the same as glyphs. The color of each path line changes continuously to highlight the particle speed.

The advection is a P-space algorithm which integrates particle velocities in the physical (longitude, latitude, meters) world coordinate system. P-space world coordinates are inverted back to their C-space logical (i, j, k) data coordinates to obtain velocities at each particle. The P-space position of each particle is integrated by its current velocity. Either a first-order (Euler) or second-order (Runge-Kutta) integration method may be used. Integration continues until the particle exits the grid or enters a singular area of the grid (i.e., land) which does not have an inverse from P-space to C-space.

The accuracy of the integration was tested by forming computer-generated velocities with each timestep pointing in a constant compass direction (E, NE, N, NW, W, SW, S, SE) in a counterclockwise order. Particles were advected at the beginning timestep and were allowed to complete several cycles. All paths were observed as purely circular. There was very little or no deviation as the particles made several passes around the circle.

The application does not have a conventional graphical user interface with menus, forms, and text fields. Instead the rendered scenery itself is the interface. For local navigation the user can travel to a point in the scene by holding the mouse at that point. For global navigation a small reference map showing the full topography can be clicked on to travel anywhere. Once in the desired area, the user can drag the mouse to change line-of-sight and to move forward or backward. Glyphs at different depths can be turned on or off by clicking a vertical legend of grid layers. Particles can be placed uniformly across a layer, or individually placed by clicking the desired location in the scene.

The user observes and controls time through a timeline at the top of the scene. The timeline covers all timesteps in the given model output. A marker in the timeline shows the current time which can be changed by clicking a different area in the line. The application maintains a timer that can be toggled on and off. When the timer is on, the application advances the current time by a given increment (default is 1 hour) and interpolates and adjusts all data and graphics to reflect the new time. Fig. 1 shows a typical Micom_Explorer scene with pathlines demonstrating the loop current in the Gulf of Mexico.

V. The Hardware

Our development hardware environment consists of a variety of SGI workstations up to and including ONYX2. The ONYX systems provide the raw power and memory for testing interactive techniques on large data sets. These systems contain up to 8 processors, 4 GB of main memory, and 64 MB of texture memory. This texture memory provides for hardware texture mapping, which outperforms software texture mapping. The Infinite Reality graphics pipe is the industry standard, for producing virtual environments. These servers come equipped with a Graphics to Video Option (GVO), which allows the capture of interactive screen sessions directly to video tape. SGI has historically transitioned their graphics technology from the large graphics servers down to their desktops. The Octane demonstrates this featuring dual processors, hardware texture mapping, and cross-bar switched memory. The INDIGO2 is our most abundant and commonly used platform. They are single-processor systems with high impact graphics and

256 MB of main memory. There are also several O2's which provide multimedia capabilities. Our target hardware is any system or workstation with a C compiler and OPENGL libraries.

VI. The Software

The operating environment at the MSRC Visualization Lab is IRIX UNIX. With the release of IRIX 6.5 we have a common version operating system between all of our SGI platforms. We are investigating, porting, and benchmarking some of these codes to other operating systems such as SOLARIS, LINUX, and NT. Our facility maintains a toolbox of commercial off-the-shelf software to assist in the display and analysis of a wide spectrum of data. These software applications range from a suite of Geographic Information Systems for analyzing geospatial data, to robust application environments such as IBM Data Explorer (DX) and Application Visualization System (AVS), to high-fidelity batch rendering applications that support a wide range of special effects such as Alias/Wavefront (now MAYA). These products provide us with a rapid prototyping environment when undertaking new projects. They provide the means to translate and understand the various data structures and how they map into 2-D and 3-D space. A full range of shareware/freeware visualization utilities available today is also supported. Eventually the application is ported to C code which exploits graphics libraries that are available on virtually every class of workstation.

VII. The People

It requires an interdisciplinary staff to accomplish these tasks. The sophistication of the visualization requires computer scientists and specialists to work closely with oceanographers, ocean modelers, geophysicists, and even warfighters to tune a visualization application to the data. Fig. 2 shows the co-author and developer of these OPENGL applications in front of a large screen projected in stereo, which provides an immersive experience.

VIII. Future Efforts

In the immediate future we want to optimize and transition our DSG surface generator to NAVOCEANO's Seafloor Data Bases Division to help them explore and quality control their gridded datasets. It may be possible to interface this application to a 2D grid editor rather painlessly. Work continues on the draping of acoustic imagery over DSG-generated 3-D surfaces. The imagery resolutions provide a challenge to create an interactive tool which allows exploration at full resolution. Dynamic paging of textures will be required, but hopefully portability can be maintained. Data correlation in general, will be exploited, but specifically, temperature

and/or salinity will be displayed with circulation in our digital exploratory environment.

IX. Acknowledgements

We would like to acknowledge the work of NAVOCEANO's Modeling and Techniques Division, particularly Dr. Phu Luong, who has subsequently left NAVOCEANO to work at the U.S. Army Corps of Engineers Waterways Experiment Station, and Dr. Eric Chassignet from the University of Miami. Both provided us with data, insight into the data, and a need to explore.

X. References

Chassignet, Eric, R. Bleck, and A. Mariano, "Data Assimilation in High Resolution Numerical Simulations of the Ocean Circulation." DoD Challenge Project Resource Request, 1999.

Horton, C., M. Clifford, D. Cole, J. Schmidt, and L. Kantha, "Operational Modeling: Semienclosed Basin Modeling at the Naval Oceanographic Office." Oceanography, Vol. 5, No. 1, 1992.

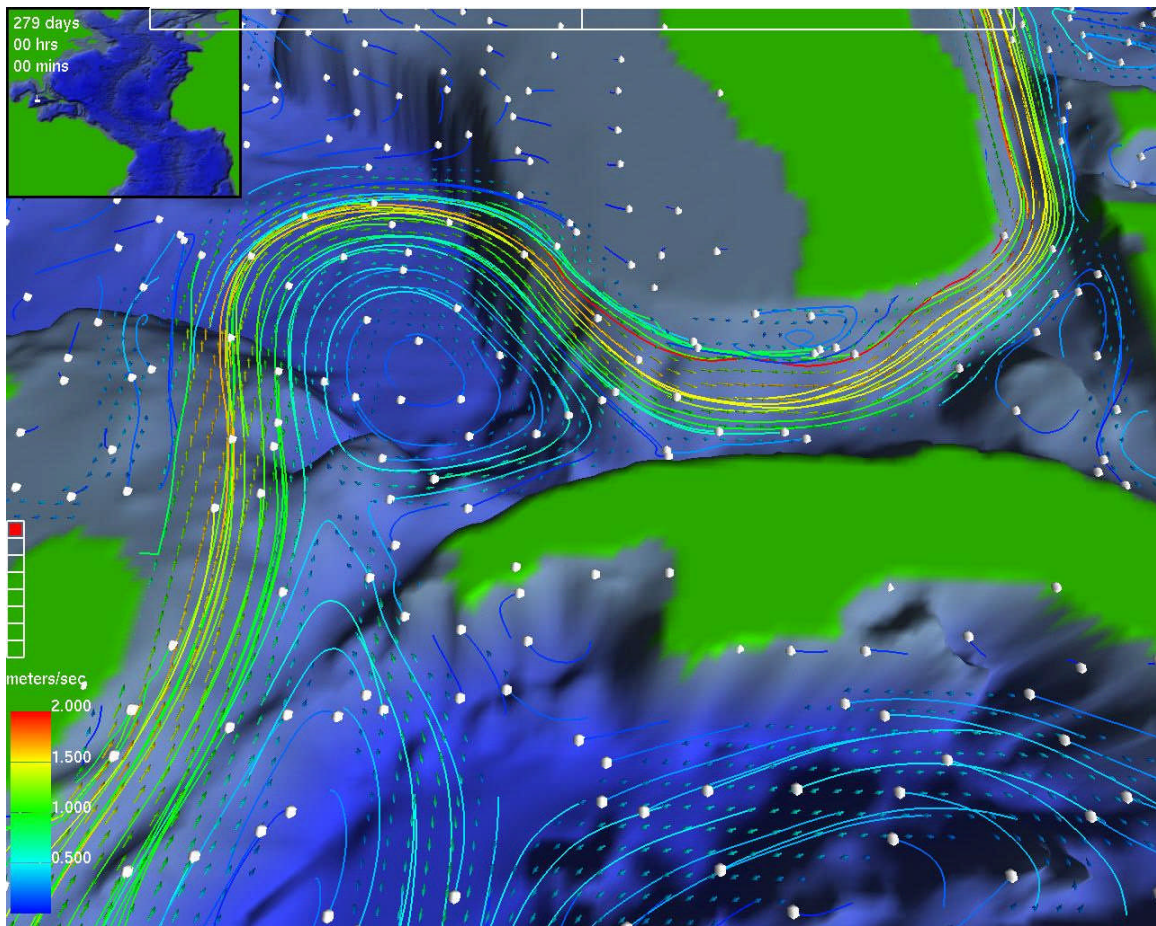


Fig. 1

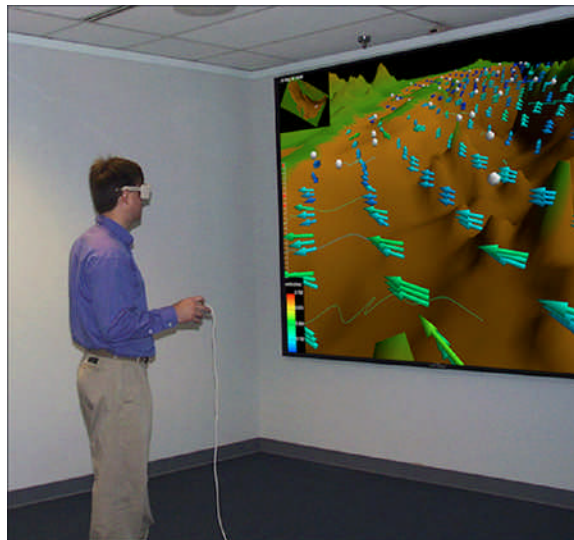


Fig. 2